

Attorney Docket No.: 59033-297872

Patent Application
Express Mail EL 971197858 US

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

TITLE OF THE INVENTION

**END-TO-END SERVICE QUALITY FOR LATENCY-INTENSIVE INTERNET PROTOCOL (IP)
APPLICATIONS IN A HETEROGENEOUS, MULTI-VENDOR ENVIRONMENT**

INVENTOR

SIDDHARTHA NAG

Prepared by

FAEGRE & BENSON LLP
3200 WELLS FARGO CENTER
1700 LINCOLN STREET
DENVER, COLORADO 80203
(303) 607-3500

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number: EL 971197858 US

Date of Deposit: November 3, 2003

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to MAIL STOP PATENT APPLICATION, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Mike DeSanctis
(Typed or printed name of person mailing paper or fee)

Mike DeSanctis
(Signature of person mailing paper or fee)

November 3, 2003
(Date signed)

**END-TO-END SERVICE QUALITY FOR LATENCY-INTENSIVE INTERNET PROTOCOL (IP)
APPLICATIONS IN A HETEROGENEOUS, MULTI-VENDOR ENVIRONMENT**

[0001] This application claims the benefit of U.S. Provisional Application No. 60/423,189, filed November 1, 2002, which is hereby incorporated by reference in its entirety. This application is a continuation-in-part of U.S. Patent Application No. 09/634,035, filed August 8, 2000, entitled “Multiplexing Several Individual Application Sessions over a Pre-Allocated Reservation Protocol Session” and U.S. Patent Application No. 10/206,402, filed July 27, 2002, entitled “Selective Encryption of Application Session Packets” both of which are hereby incorporated by reference in their entirety.

COPYRIGHT NOTICE

[0002] Contained herein is material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction of the patent disclosure by any person as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all rights to the copyright whatsoever.

BACKGROUND

Field

[0003] Embodiments of the present invention generally relate to the delivery of true end-to-end application Quality of Service (QoS) over Internet Protocol (IP) networks. More particularly, embodiments of the present invention relate to techniques for pre-allocating an aggregated end-to-end network reservation protocol session between heterogeneous, multi-vendor environments and thereafter sharing the reservation protocol session among multiple individual application sessions, such as voice, video, and real-time media applications, by multiplexing the multiple individual application flows running across the end-to-end network thereon to achieve desired QoS needs.

Description of the Related Art

[0004] The consolidation and transfer of voice and voice-band data (e.g., fax and analog modems) with data services over public packet networks, such as the Internet, is rapidly gaining acceptance. However, significant work remains to support the high-availability and tight quality of service (QoS) requirements needed to support voice, video, and real-time content applications over IP networks.

[0005] A variety of IP QoS mechanisms are currently available. Some of the more prevalent examples include IP router techniques such as localized queuing and prioritized packet classifications as well as standardized networking protocols such as Resource Reservation Protocol (RSVP), Differentiated Services (DiffServ), and Multiprotocol Label Switching (MPLS). While these solutions have been heavily marketed as QoS solutions, they have met with poor customer acceptance because they represent, at best, only partial solutions for IP QoS. Even when used in conjunction with one another, there are several major deficiencies with the current approaches.

[0006] First, they are static in nature and non-adaptive to real-time network load and delay conditions which have critical effects on application performance such as voice. Additionally, since these solutions typically require a manual and pre-determined traffic engineering process to identify optimal network routing and bandwidth, once the design is implemented in the network there is no ability to make intelligent routing and control decisions to compensate for dynamic network behavior.

[0007] Second, today's solutions are transport-centric with no awareness of the individual application flows running across the end-to-end network. Thus, since no distinction can be made between packets, such as a voice packet and an ordinary data packet, for example, there is once again no ability to make intelligent routing and control decisions to ensure end-to-end application QoS.

[0008] Finally, there are a variety of other deficiencies depending on the mechanism being used. These include "per-router-hop behaviors" with no tightly-coupled, end-to-end

QoS integration, high complexity and overhead; and extensive requirements for ongoing traffic engineering and network design as mentioned above.

[0009] In summary, today's IP QoS mechanisms and protocols provide a part of the solution to the problem, but have proven to be unworkable in real-world, voice, video and data IP networks.

SUMMARY

[00010] Apparatus and methods are described for delivering end-to-end application Quality of Service (QoS) over Internet Protocol (IP) networks. According to one embodiment, According to one embodiment, a portion of available bandwidth between a first and second network device is reserved as a Quality of Service (QoS) resource pool for real-time communication sessions among users of a first and second user community. The first network device is communicatively coupled with a packet network and associated with the first user community. The second network device is communicatively coupled with the packet network and associated with the second user community. End-to-end application QoS is provided between the first and second user communities by selectively admitting real-time communication sessions between the first user community and the second user community based upon currently available resources associated with the QoS resource pool and multiplexing the real-time communication sessions over a reservation protocol session between the first and second network devices.

[00011] Other features of embodiments of the present invention will be apparent from the accompanying drawings and from the detailed description that follows.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[00012] Embodiments of the present invention are illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[00013] **Figure 1A** illustrates a simplified logical architecture of an enterprise VoIP network according to one embodiment of the present invention.

[00014] **Figure 1B** conceptually illustrates interactions between two media aggregation managers according to one embodiment of the present invention.

[00015] **Figure 2** is an example of a network device in which one embodiment of the present invention may be implemented.

[00016] **Figure 3** is a high-level block diagram of a media aggregation manager according to one embodiment of the present invention.

[00017] **Figure 4** is a simplified, high-level flow diagram illustrating application session processing according to one embodiment of the present invention.

[00018] **Figure 5** is a simplified, high-level flow diagram illustrating application session establishment processing according to one embodiment of the present invention.

[00019] **Figure 6** illustrates interactions among local and remote media aggregation manager functional units according to one embodiment of the present invention.

[00020] **Figure 7** is a flow diagram illustrating Registration, Admission, Status (RAS) signaling processing according to one embodiment of the present invention.

[00021] **Figure 8** is a flow diagram illustrating call signaling processing according to one embodiment of the present invention.

[00022] **Figure 9** is a flow diagram illustrating control signaling processing according to one embodiment of the present invention.

[00023] **Figure 10** is a flow diagram illustrating media/control transmission processing according to one embodiment of the present invention.

[00024] **Figure 11** is a flow diagram illustrating media/control reception processing according to one embodiment of the present invention.

[00025] **Figure 12** conceptually illustrates application session establishment in an H.323 environment according to one embodiment of the present invention.

[00026] **Figure 13** conceptually illustrates H.323 signaling and media flow according to an embodiment in which call management is performed external to the media aggregation managers.

[00027] **Figure 14A** illustrates the encapsulated (“MUX”) packet format according to one embodiment of the present invention in which address replacement is performed by the LMAM.

[00028] **Figure 14B** illustrates media transmission in both directions according to the encapsulated packet format of **Figure 14A**.

[00029] **Figure 15A** illustrates the encapsulated (“MUX”) packet format according to another embodiment of the present invention in which address replacement is performed by the RMAM.

[00030] **Figure 15B** illustrates media transmission in both directions according to the encapsulated packet format of **Figure 15A**.

DETAILED DESCRIPTION

[00031] Apparatus and methods are described for providing end-to-end application quality of service (QoS) over IP networks. Embodiments of the present invention seek to provide a scalable and flexible architecture that enables efficient provisioning of reserved bandwidth to multiple application flows running across the end-to-end network by multiplexing the individual application flows over a pre-allocated reservation protocol session thereby providing true end-to-end QoS support. The pre-allocated reservation protocol session may take into consideration current network resources and estimated usage of network resources, such as bandwidth, based upon historical data. For example, the amount of pre-allocated resources may vary due to different loads being offered at different times of day and/or day of week. Additionally, the pre-allocated reservation protocol session may be dynamically adjusted to account for actual usage that surpasses the estimated usage or actual usage that falls below the estimated usage.

[00032] According to one embodiment, a more intelligent approach is employed in connection with initiation and maintenance of a large number of reservations. Rather than establishing and maintaining a reservation protocol session for each application flow that requires real-time response, which results in many independent reservation protocol sessions and high overhead, a single reservation protocol session may be pre-allocated and subsequently dynamically shared among the application flows by aggregating the associated media packets and transmitting them as a multiplexed media stream over the end-to-end IP virtual connection QoS pipe. As a result, in this embodiment, a single, dynamic end-to-end QoS pipe may be maintained between two different user communities using a pre-allocated RSVP session between a pair of media aggregation managers. The media aggregation managers multiplex (and optionally selectively encrypt) outbound voice packets onto the pre-allocated RSVP session (and decrypt) and demultiplex inbound voice packet received over the pre-allocated RSVP session, thereby sharing a common RSVP session and reducing the computational resources required by the network to provide real-time response for multiple application flows. Advantageously, in this manner, application-level intelligence, dynamic

traffic engineering, and IP route control capabilities are extended to existing IP QoS protocols thereby making it feasible to use reservation protocols, such as RSVP, for large numbers of applications that require real-time performance and QoS, such as VoIP services.

[00033] In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form. Embodiments of the present invention include various steps, which will be described below. The steps may be performed by hardware components or may be embodied in machine-executable instructions, which may be used to cause a general-purpose or special-purpose processor programmed with the instructions to perform the steps. Alternatively, the steps may be performed by a combination of hardware and software.

[00034] Embodiments of the present invention may be provided as a computer program product which may include a machine-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform a process. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnetic or optical cards, flash memory, or other type of media / machine-readable medium suitable for storing electronic instructions. Moreover, the present invention may also be downloaded as a computer program product, wherein the program may be transferred from a remote computer to a requesting computer by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection).

[00035] While, for convenience, embodiments of the present invention are described with reference to particular existing signaling, control, and communications protocol standards, such as International Telecommunication Union Telecommunication Standardization Section (ITU-T) Recommendation H.225.0 entitled "Call Signalling

Protocols and Media Stream Packetization for Packet-based Multimedia Communication Systems," published February 1998 (hereinafter H.225.0); ITU-T Recommendation H.245 entitled "Control Protocol for Multimedia Communication," published May 1999 (hereinafter H.245); ITU-T Recommendation H.323 entitled "Packet-based Multimedia Communications Systems," published September 1999 (hereinafter H.323); and a particular bandwidth reservation protocol (i.e., RSVP), the present invention is equally applicable to various other signaling, control, communications and reservation protocols. For example, Session Initiation Protocol (SIP) may be employed to create, modify, and terminate application sessions with one or more participants. SIP is described in M. Handley et al., "SIP: Session Initiation Protocol," RFC 2543, Network Working Group, March 1999, which is hereby incorporated by reference. Furthermore, it is contemplated that embodiments of the present invention will be applicable to various proprietary signaling and media transport protocols such as those employed between and among IP Private Branch Exchange (PBX) systems and IP phones of various vendors.

[00036] In addition, for sake of brevity, embodiments of the present invention are described with reference to a specific application (i.e., VoIP) in which individual flows may be multiplexed over a pre-allocated bandwidth reservation protocol session. Nevertheless, embodiments of the present invention are equally applicable to various other network applications or services that are latency intensive (e.g., affected by jitter and/or transmission delays) and/or that require real-time performance, such as applications based on human interactions (e.g., collaborative software, online/Web collaboration, voice conferencing, and video conferencing), and real-time data communication and/or exchange, such as market data applications, financial transactions, and the like.

Terminology

[00037] Brief definitions of terms used throughout this application are given below.

[00038] A "media aggregation manager" may generally be thought of as a network device, such as an edge device at the ingress/egress edges of a user community or enterprise

site or a group of one or more software processes running on one or more servers that provides application/protocol specific multiplexing/demultiplexing of media traffic through a pre-allocated reservation protocol session also referred to as an IP virtual connection (VC) QoS pipe. In various embodiments, the media aggregation manager may include more or less functionality depending upon the target usage environment. For example, in one embodiment, the media aggregation manager implements signaling protocol functionality to communicate with IP phone sets, terminals, and/or other IP telephony products. In another embodiment, in which the media aggregation manager is intended to cooperate with an IP PBX system, the media aggregation manager acts as a signaling and/or media gateway and an H.323 gatekeeper.

[00039] A “reservation protocol” generally refers to a protocol that may be employed to communicate information regarding a desired level of service for a particular application flow. An example of an existing bandwidth reservation protocol is RSVP.

[00040] A “user community” generally refers to a group of users residing on a common network at a given location. For example, employees on an enterprise network at a given location, users of a particular Internet service provider (ISP) at a given location, subscribers to a particular long distance carrier in a given region, or other users accessing a distributed IP network via a common access point may represent a user community.

[00041] A “reservation protocol session” generally refers to a set of reserved network resources established and maintained between two or more network devices that serve as proxies for application endpoints residing behind the proxies. An example, of a reservation protocol session is an RSVP session between two media aggregation managers.

[00042] An “application session” generally refers to a session established and maintained between two or more terminals. According to embodiments of the present invention, one or more application sessions may be multiplexed onto a single reservation protocol session thereby reducing the overhead for establishing and maintaining multiple reservation protocol sessions.

[00043] A “terminal” generally refers to a LAN-based endpoint for media transmission, such as voice and/or voice-based data transmission. Terminals may be capable of executing one or more networked applications programs. Examples of terminals include IP phones and computer systems running an Internet telephony application, such as CoolTalk or NetMeeting.

[00044] A “tunnel” generally refers to a logical transmission medium through which packets of one protocol encapsulated or wrapped in a packet of another protocol are transmitted via the protocol of the wrapper. According to one embodiment, voice and/or voice-band data packets are encrypted proximate to the source for secure transmission over one or more public internetworks, such as the Internet, and then decrypted proximate to the destination.

[00045] An “application” or “endpoint” generally refers to a software program that is designed to assist in the performance of a specific task, such as Internet telephony, online collaboration, video conferencing, or exchange of mission critical data.

[00046] An “application flow” generally refers to the data associated with an application session. An example of an application flow is a media stream, such as a continuous sequence of packetized voice and/or voice-band data transmitted over a network.

[00047] A “tag,” in the context of the described embodiment, generally refers to information that is appended to application generated packets, such as Real-time Transport Protocol (RTP) packets or Real-time Transport Control Protocol (RTCP) packets, that allows the proxy endpoints of the reservation protocol session to transmit encapsulated packets to the appropriate remote application/endpoint (RA). According to one embodiment of the present invention, a tag includes address information, such as the destination network address of the terminal upon which the destination application/endpoint resides. When a media aggregation manager is employed in connection with a transport protocol and control protocol (such as RTP and RTCP) that use different channels or ports for control and data, control and data packets may be multiplexed onto the reservation protocol session as well by including protocol dependent control information. Then, the remote media aggregation manager may

strip the tag from the encapsulated packet and determine the appropriate channel/port of the remote application/endpoint on which to forward the resulting packet based upon the additional protocol dependent control information within the tag. Advantageously, in this manner, two layers of multiplexing may be achieved; (1) a first layer that allows identification of the appropriate application at the remote media aggregation manager; and (2) a second layer that specifies a subclass/subprocess within an application.

Media Aggregation Overview

[00048] The architecture described herein seeks to resolve scalability problems observed in current reservation protocols. These scalability issues have slowed the adoption of reservation protocols in network environments where multiple applications must be provided with certainty regarding a minimum reserved bandwidth. The architecture described herein additionally seeks to address scalability problems associated with current security solutions for IP telephony product offerings. Furthermore, embodiments of the present invention seek to extend application-level intelligence, dynamic traffic engineering, and IP route control capabilities to existing IP QoS protocols.

[00049] **Figure 1A** illustrates a simplified logical architecture of an enterprise VoIP network according to one embodiment of the present invention. In the example application environment depicted, two enterprise locations (e.g., branch offices or campuses) 110 and 120 may use a distributed IP network 100, such as the Internet, as a transmission medium for the transfer of voice and voice-band data. Because vendors of hardware and software for IP telephony products typically employ proprietary signaling protocols, customers are required to implement homogeneous VoIP environments with IP PBX systems and IP phone sets from the same vendor. However, according to one embodiment of the present invention, signaling and/or media gateway functionality is provided in an intermediate device, e.g., media aggregation managers 115 and 125 associated with each enterprise location 110 and 120, respectively, that allows interoperation among heterogeneous VoIP environments. For example, as a result of the protocol bridging functionality described herein, a particular

enterprise location employing one vendor's IP telephony solution and products may seamlessly communicate with another enterprise location employing another vendor's IP telephony solution and products. Additionally, the media stream that flows between the enterprise locations provides very high QoS across heterogeneous, multi-vendor environments. Furthermore, embodiments of the present invention allow enterprise locations to use IP phone sets of one vendor and an IP PBX system of another.

[00050] Returning to **Figure 1A**, enterprise location 110 includes one or more IP phone sets or terminals 111 – 113 and a media aggregation manager 115. Similarly, enterprise location 120 includes one or more IP phone sets or terminals 121 – 123 and a media aggregation manager 125. The enterprise locations are communicatively coupled via an IP network 100. In this example, two different IP PBX call management agents 130 and 150 are depicted. The IP PBX call management agents represent existing or future call signaling functionality, such as that provided by Cisco Systems, Inc.'s Integrated Communications System (ICS) 7750 with CallManager software, 3Com Communications' NBX® 100 Communications System with 3Com® Superstack® NBX Call Processor software, Avaya Inc.'s IP600 Communications Server with MultiVantage™ software, Siemens' SURPASS™ NetManager™, NEC's CX6100-CA call agent, and/or Alcatel's OmniPCX 4400 Call Server. At any rate, one IP PBX call management agent 130 is associated with enterprise location 110 and the other IP PBX call management agent 150 is associated with enterprise location 120. Advantageously, by bypassing the Public Switched Telephone Network (PSTN) (at least for long distance service) such a VoIP environment allows users to make voice calls anywhere in the world as part of their typically fixed-price Internet access rate. In any event, The IP PBX call management agents 130 and 150 may be from the same or different vendors and the IP phone sets 111 – 113 and 121 – 123 may be from the same or different vendors as the associated IP PBX call management agents 130 and 150, respectively.

[00051] **Figure 1B** conceptually illustrates interactions between two media aggregation managers 115 and 125 according to one embodiment of the present invention. According to

one embodiment, the media aggregation managers 115 and 125 act as reservation protocol proxies on behalf of the terminals 111, 112, 121, and 122. For example, the media aggregation managers 115 and 125 establish and maintain a reservation session, such as an RSVP session, between each other by exchanging reservation signaling messages 160. Subsequently, rather than establishing additional reservation protocol sessions, the media aggregation managers 115 and 125 respond to reservation requests from the terminals 111, 112, 121, and 122 by dynamically allocating the reserved resources, such as bandwidth, associated with the reservation protocol session to corresponding application sessions. In this manner, multiple application sessions may share the reservation session by multiplexing media packets onto the reservation session as described further below.

[00052] According to one embodiment, the media aggregation managers 115 and 125 may additionally act as tunnel endpoints through which encrypted voice and/or voice-band data and exchanged among the terminals 111, 112, 121, and 122. For example, one or more tunnels may be established between the media aggregation managers 115 and 125 through the pre-allocated reservation session by exchanging tunnel protocol signaling messages 165.

[00053] In this example, an IPVC QoS pipe 170 is established using admission control signaling messages 180. The media/control stream 171 is carried over the pre-allocated reservation session between media aggregation manager 115 and media aggregation manager 125. A multiplexed media/control stream represents one way to handle certain transport and control protocol combinations, such as RTP and RTCP, that use different channels or ports for control and data. In alternative embodiments, the reservation protocol session 170 may not need to distinguish between control and data.

[00054] While in some embodiments described herein, the media aggregation managers 115 and 125 are discussed as if they are autonomous network edge devices, it should be kept in mind that according to various other embodiments of the present invention some or all of the functionality of a media aggregation manager might be integrated with existing network devices, such as bridges, routers, switches, gateways, servers, and the like. Additionally, while only a single aggregated reservation protocol session between two media

aggregation managers 115 and 125 is described in connection with the present example, it should be appreciated that each media aggregation manager 115 and 125 may support multiple, heterogeneous reservation protocol sessions capable of providing heterogeneous application flows among multiple user communities. According to one embodiment of the present invention, regardless of the number of terminals or application/endpoints, application flows may be provided with reserved bandwidth between any and all pairs of terminals of N user communities by establishing and sharing no more than N^2 reservation protocol sessions.

Network Device Overview

[00055] An exemplary machine in the form of a network device 200, representing an exemplary media aggregation manager 115, in which features of the present invention may be implemented will now be described with reference to **Figure 2**. In this simplified example, the network device 200 comprises a bus or other communication means 201 for communicating information, and a processing means such as one or more processors 202 coupled with bus 201 for processing information. Networking device 200 further comprises a random access memory (RAM) or other dynamic storage device 204 (referred to as main memory), coupled to bus 201 for storing information and instructions to be executed by processor(s) 202. Main memory 204 also may be used for storing temporary variables or other intermediate information during execution of instructions by processor(s) 202. Network device 200 also comprises a read only memory (ROM) and/or other static storage device 206 coupled to bus 201 for storing static information and instructions for processor 202. Optionally, a data storage device (not shown), such as a magnetic disk or optical disc and its corresponding drive, may also be coupled to bus 201 for storing information and instructions.

[00056] One or more communication ports 225 may also be coupled to bus 201 for allowing various local terminals, remote terminals and/or other network devices to exchange information with the network device 200 by way of a Local Area Network (LAN), Wide Area Network (WAN), Metropolitan Area Network (MAN), the Internet, or the public switched telephone network (PSTN), for example. The communication ports 225 may include various

combinations of well-known interfaces, such as one or more modems to provide dial up capability, one or more 10/100 Ethernet ports, one or more Gigabit Ethernet ports (fiber and/or copper), or other well-known interfaces, such as Asynchronous Transfer Mode (ATM) ports and other interfaces commonly used in existing LAN, WAN, MAN network environments. In any event, in this manner, the network device 200 may be coupled to a number of other network devices, clients and/or servers via a conventional network infrastructure, such as a company's Intranet and/or the Internet, for example.

Media Aggregation Manager

[00057] **Figure 3** is a high-level block diagram of a media aggregation manager according to one embodiment of the present invention. By interconnecting a plurality of distributed media aggregation managers, such as media aggregation manager 300, a media and signaling gateway architecture is provided for moving media from a potentially multi-vendor PBX community through a high quality pipe to one or more other potentially multi-vendor PBX communities. For example, several application flows (e.g., VoIP calls) may be selectively encrypted and multiplexed over a pre-allocated reservation protocol session, such as a pre-allocated RSVP pipe, e.g., IPVC QoS pipe 170. Advantageously, the multiplexing of application flows reduces the computational resources required by the network to provide reserved bandwidth, e.g., guaranteed bandwidth, for multiple application flows. Additionally, the selective encryption at the media aggregation manager 300 is more elegant and scalable than RBE solutions due in part to the ability of the media aggregation manager's ability to maintain state on a packets-to-call mapping basis. Furthermore, in a VoIP environment, the logical positioning of the media aggregation managers 300 relative to the terminals enables it to (1) apply application-level intelligence to determine which packets can traverse the IPVC QoS pipe 170, (2) gather and maintain routing information and real-time route performance data to ensure efficient bandwidth utilization, least-cost routing, and optimal service performance, and (3) perform selective encryption (e.g., optional encryption only on voice

related packets) rather than encrypting every packet that hits the router as would be the case in a RBE solution.

[00058] In the example depicted, the source media aggregation manager receives media packets from its local terminals and transmits encrypted multiplexed media to the destination aggregation manager. The destination aggregation manager receives the encrypted multiplexed media and routes media packets to the appropriate terminal(s) of its local terminals by performing demultiplexing and decryption.

[00059] In this example, the media aggregation manger 300 includes an application/protocol specific media multiplexor 350, an application/protocol specific media demultiplexor 360, a media encryptor 355, a media decryptor 365, an admission control manager 315, a tunneling control manager 320, a signaling gateway 330, a generic resource manager 340, and a resource pool 345. In a software implementation, instances of the media multiplexor 350, media demultiplexor 360, and admission control manager 315 may be created for each particular application/protocol needed to allow communications between terminals of the geographically diverse user communities. Similarly, appropriate instances of the signaling gateway 330 can be instantiated depending upon signaling translations required by the particular environment. Importantly, it should be appreciated that the particular partitioning of functionality described with reference to this example is merely illustrative of one or many possible allocations of functionality.

[00060] According to the embodiment depicted, the resource manager 340 establishes and maintains one or more pre-allocated reservation protocol sessions between the local media aggregation manager and one or more remote media aggregation managers. The resource manager 340 optionally interfaces with a centralized entity that provides information relating to the characteristics and estimated amount of resources for the pre-allocated reservation protocol sessions. Alternatively, a network administrator may provide information to the resource manager 340 relating to desired characteristics of the pre-allocated reservation protocol sessions. The resource manager 340 also tracks active

application sessions for each reservation protocol session and the current availability of resources for each reservation protocol session in the resource pool 345.

[00061] The signaling/media gateway 330 interfaces with the local terminals and one or more other remote signaling gateways (RSGs) associated with other user communities to perform signaling/media translation between the potentially different proprietary signaling protocols employed at the local and remote user communities. Additionally, according to one embodiment, the signaling gateway 330 may be logically interposed between the local terminals and the associated call management agent to perform local signaling and/or media translation between local terminals and the associated call management agent. In this manner, IP phones and IP PBXs of different vendors can be mixed and matched to allow media from one potentially multi-vendor PBX community to be moved through a high quality pipe to one or more other potentially multi-vendor PBX communities.

[00062] The tunneling control manager 320 interfaces with the media encryptor 350 and one or more other remote tunneling control managers (RTCMs) associated with other user communities to agree upon encryption, such as Message Digest 5 (MD5), RSA Data Encryption Standard (DES) or other encryption standard, key management, and/or a tunneling protocol to be employed for a particular application session, such as existing or future versions of the IP Security (IPSec) Protocol, generic routing encapsulation (GRE), Layer 2 Forwarding (L2F), Layer 2 Tunneling Protocol (L2TP), or the Point-to-Point Tunneling Protocol (PPTP).

[00063] The media encryptor 355 receives media packets from the local terminals (not shown) and selectively encrypts the media packets for exchange with the media decryptor 365 of the remote media aggregation manager as previously agreed upon by the participating tunneling control managers based upon the application session with which the media packets are associated. In this manner, security may be configured on an application session basis (e.g., a call-by-call basis).

[00064] The admission control manager 315 interfaces with local terminals (not shown) associated with a particular user community, the media multiplexor 350, the resource

manager 340, and one or more other remote media aggregation managers associated with other user communities. Importantly, in one embodiment, the media multiplexor 350 hides the details of how reserved resources are internally allocated and managed, thereby allowing the local terminals to use existing reservation protocols, such as RSVP, without change.

[00065] The media multiplexor 350 receives selectively encrypted media packets from the media encryptor 355 and appropriately translates/encapsulates the packets for communication with the media demultiplexor 360 of the remote media aggregation manager in accordance with the aggregation technique described further below. When application flows are established and terminated, the admission control manager 315 interfaces with the resource manager 340 to allocate and deallocate resources, respectively.

[00066] The media demultiplexor 360 interfaces with the media decryptor 365 to supply the media decryptor 365 with the selectively encrypted media packets by demultiplexing the respective application flows from the pre-allocated reservation protocol session. The media decryptor 365 then decrypts the media packets, if necessary, and forwards them to the appropriate local terminals (not shown).

[00067] The admission control manager 315 exchanges admission control signaling messages with remote admission control managers and configures the local application/endpoint (LA) to send media to transmitted to the local media aggregation manager after an application session has been established with a remote media aggregation manager. For VoIP using the H.323 protocol, the admission control manager 315 may include RAS, call control, and call signaling processing.

[00068] When application flows are established and terminated, the admission control manager 315 interfaces with the resource manager 340 to allocate and deallocate resources, respectively.

[00069] In operation, two resource managers cooperate to establish a pre-allocated reservation protocol session between a local media aggregation manager (LMAM) and a remote media aggregation manager (RMAM). The resource managers make a reservation that is large enough to accommodate the anticipated load offered by applications that need to

communicate over the reservation protocol session. Subsequently, a local media encryptor (LME) associated with the LMAM provides admission control for application flows between one or more terminals of the LMAM and the RMAM with the assistance of the local and remote admission control managers and the local and remote resource managers. If sufficient resources, such as bandwidth, are available over the pre-allocated reservation protocol session, then the LME selectively encrypts the application flows and the local media multiplexor (LMM) multiplexes the application flows for transmission over the pre-allocated reservation protocol session. On the receiving end, the remote media demultiplexor (RMDX) demultiplexes the application flows and sends them to their intended destinations through the remote media decryptor (RMD) which performs any necessary decryption. The typical admission control manager 315 will be a player in the path of the application protocol for setting up the connection between two or more application endpoints; hence, it may be instrumented to modify the path of the media packets to flow through the LME, LMM, the remote media encryptor (RME), and the remote media multiplexor (RMM).

[00070] In brief, after an application session has been associated with the pre-allocated reservation protocol session, the application/endpoints may use a transport protocol and/or a control protocol, such as RTP and/or RTCP to exchange encrypted media packets between them. The media packets may carry various types of real-time data, such as voice, voice-band data, video, multi-media, real-time market data, mission critical data, or other data for human interactions or collaboration. Media packets from a data source are optionally encrypted by the local media encryptor 355, tagged by the local media multiplexor 350, and sent over the reserved path to one or more media demultiplexors 360 corresponding to the data destination. As illustrated below, the media demultiplexor 360 strips the tag before the media packets are forwarded, the media decryptor 360 performs decryption processing, and then the tag information is used to determine the ultimate destination of the data packet.

[00071] According to one embodiment, from the perspective of the local terminals, they are establishing and using reservation protocol sessions for each application flow and communicating in the clear. However, in reality, the media aggregation manger 300 shares

the pre-allocated reservation protocol session among multiple application flows and transparently performs encryption and/or decryption as necessary.

[00072] As will be described further below, a specific example of the use of this architecture is in connection with the use of the H.323 protocol for VoIP calls. Typically, an H.323 Gatekeeper is used by endpoints to help in address resolution, admission control etc. So, for the H.323 protocol, the gatekeeper is a convenient place for the media multiplexor 350 and/or media encryptor 355 to reside. Alternatively, the media aggregation manager 300 may implement the H.323 gatekeeper functionality and act as a gatekeeper for devices, such as IP PBXs.

[00073] Note that in this description, in order to facilitate explanation, the media aggregation manager 300 is generally discussed as if it is a single, independent network device or part of single network device. However, it is contemplated that the media aggregation manager 300 may actually comprise multiple physical and/or logical devices connected in a distributed architecture; and the various functions performed may actually be distributed among multiple network devices. Additionally, in alternative embodiments, the functions performed by the media aggregation manager 300 may be consolidated and/or distributed differently than as described. For example, any function can be implemented on any number of machines or on a single machine. Also, any process may be divided across multiple machines. Specifically, the media multiplexor 350 and the media encryptor 355 may be combined as a single functional unit or the multiplexing and encrypting processing may be performed in the opposite order than described above. Similarly, the media demultiplexor 360 and the media decryptor 365 may be combined as a single functional unit or the demultiplexing and decrypting processing may be performed in the opposite order than described above, e.g., encryption may be performed before or after multiplexing. Finally, encryption may be performed at various levels of the application flow. For example, encryption may be performed on the media and/or control information directly, the media and/or control packets, or on multiplexed media and/or control packets.

Sharing a Pre-Allocated Reservation Protocol Session

[00074] **Figure 4** is a simplified, high-level flow diagram illustrating application session processing according to one embodiment of the present invention. In one embodiment, the processing blocks described below may be performed under the control of a programmed processor, such as processor 202. However, in alternative embodiments, the processing blocks may be fully or partially implemented by any programmable or hard-coded logic, such as Field Programmable Gate Arrays (FPGAs), TTL logic, or Application Specific Integrated Circuits (ASICs), for example.

[00075] In this example, it is assumed that a reservation protocol session has been previously established. The pre-allocated reservation protocol session preferably takes into consideration current network resources and estimated usage of network resources, such as bandwidth, based upon historical data. For example, the amount of pre-allocated resources may vary due to different loads being offered at different times of day and/or day of week.

[00076] At any rate, at decision block 410, the media aggregation manager 300 determines the type of event that has occurred. If the event represents the receipt of an application session establishment request from a local terminal, then processing proceeds to decision block 420. If the event represents the receipt of media packets from a local application/endpoint, then processing continues with decision block 450. If the event represents the receipt of a media packet from a remote application/endpoint, then control passes to processing block 460. If the event represents the receipt of an application session termination request, then processing continues with processing block 470.

[00077] At decision block 420, a determination is made whether resources are available to meet the needs identified in the application session establishment request. For example, the resource manager 340 may determine if sufficient bandwidth is available on an appropriate pre-allocated reservation protocol session by comparing a minimum bandwidth specified in the application session establishment request to a bandwidth availability indication provided by the resource pool 345.

[00078] If adequate resources are available to provide the requestor with the minimum resources requested, processing continues with processing block 430 where application session establishment processing is performed. Application session establishment processing is described below with reference to **Figure 5**. Otherwise, if there are insufficient resources to accommodate the application session establishment request, processing branches to processing block 440. At processing block 440, the media aggregation manager 300 may reject the application session establishment request. Alternatively, the media aggregation manager 300 may continue the application session establishment process and provide a best effort service for the request (without the use of pre-allocated resources, e.g., the IPVC QoS pipe 170).

[00079] At processing block 450, media packets received from a local application/endpoint are selectively encrypted depending upon the application session with which they are associated, tagged, and sent over the network to the destination using the previously reserved resources (e.g., the pre-allocated reservation protocol session 170). The tagging and multiplexing of media packets onto the pre-allocated reservation protocol session will be discussed in detail below.

[00080] At processing block 460, potentially encrypted and multiplexed media packets received from a remote application/endpoint are decrypted, if necessary, and forwarded to the appropriate local application/endpoint. For example, the packets may be sent to the appropriate local application/endpoint based upon an examination of the tag information added by the remote media aggregation manager.

[00081] At processing block 470, in response to an application session termination request, resources allocated to this application session are relinquished and made available for other application sessions. For example, the resource manager 340 may update an indication of available resources in the resource pool 345 for the pre-allocated reservation protocol session associated with the terminated application session.

[00082] **Figure 5** is a simplified, high-level flow diagram illustrating application session establishment processing according to one embodiment of the present invention. In

the present example, application session establishment processing begins with processing block 510. At processing block 510, the requested resources are allocated to the application session. According to one embodiment, the local resource manager 340 creates a new application session entry, in the resource pool 345, containing an indication of the resources granted to the application session.

[00083] At decision block 520, a determination is made whether the desired remote application/endpoint is available to participate in the application session. If so, processing proceeds to processing block 530; otherwise, processing branches to processing block 560.

[00084] Assuming the desired remote application/endpoint is available to participate in the application session, then at processing block 530, the local application/endpoint and the remote application/endpoint are configured to send media packets associated with the application session to the local and remote media multiplexors, respectively.

[00085] At processing block 540, the local and remote media multiplexors and demultiplexors are configured in accordance with the application session. For example, as described further below, a lookup table may be maintained by the media multiplexor 350 or media demultiplexor 360 to translate the source network address of the local application/endpoint to the destination network address of the remote application/endpoint.

[00086] Finally, at processing block 550, the local and remote media encryptors 355 and decryptors 365 are optionally configured in accordance with the desired level of security for the application session. Exemplary standards-based encryption options include Message Digest 5 (MD5), RSA Data Encryption Standard (DES), and Triple DES encryption. Exemplary tunneling options include the IP Security (IPSec) Protocol, generic routing encapsulation (GRE), Layer 2 Forwarding (L2F), Layer 2 Tunneling Protocol (L2TP), or the Point-to-Point Tunneling Protocol (PPTP).

[00087] **Figure 6** illustrates interactions among local and remote media aggregation manager functional units according to one embodiment of the present invention. In general, according to the present example, the media aggregation managers abstract the true application session endpoints from each other and serve as proxies for their respective local

applications/endpoints. The media aggregation managers accomplish this by intercepting messages originating from their respective local applications/endpoints and modifying the messages to make themselves appear as the actual application flow originators/recipients.

[00088] In this example, for simplicity, it is assumed that a single local application/endpoint (LA) is establishing an application session with a single remote application/endpoint (RA) over a pre-allocated reservation protocol session 690 between a local media aggregation manager (LMAM) logically associated with or geographically proximate to the LA and a remote media aggregation manager (RMAM) logically associated with or geographically proximate to the RA.

[00089] The LA transmits a request to connect to the RA to the LMAM (670). The LACM inquires of the local resource manager (LRM) whether sufficient resources are currently available to accommodate the LA's request (672). The LRM indicates the availability or inavailability of available resources to the LACM (674).

[00090] Assuming sufficient resources are available to provide the reserved resources the LA needs for the requested connection to the RA, then the LACM asks the RACM if the RA is available (676). In response to the LACM's request, the RACM queries the RA to determine its present availability (678). The RA indicates whether or not it is currently available to participate in an application session (680).

[00091] Assuming, the RA indicates that it is available, then the RACM communicates the RA's availability to the LACM (682). In response to the availability of the RA, the LACM directs the RACM to proceed with establishment of a connection between the LA and RA.

[00092] Having determined that a connection is feasible, the LACM and RACM proceed to configure their media multiplexors and media demultiplexors for the LA-RA connection. The LACM configures the local media multiplexor (LMM) to tag media originated from the LA for routing to the RA and to send the resulting encapsulated media packets to the remote media demultiplexor (RMDX) (686). The LACM further configures

the local media demultiplexor (LMDX) to forward media packets that are received from the RMM and tagged as being associated with the LA-RA connection to the LA (690).

[00093] Similarly, the RACM configures the remote media demultiplexor (RMDX) to forward media packets that are received from the LMM and tagged as being associated with the LA-RA connection to the RA (688). The RACM also configures the remote media multiplexor (RMM) to tag media originated from the RA for routing to the LA and to send the resulting encapsulated media packets to the local media demultiplexor (LMDX) (692).

[00094] Once the media multiplexors and media demultiplexors have been appropriately configured for the LA-RA connection, the LACM and the RACM inform their application/endpoints to commence transmission of media to the LME and the RME, respectively 694 and 696. Thus, the media aggregation managers appear to their respective application/endpoints as the actual application flow originators/recipients and subsequently serve as proxies for their respective application/endpoints.

[00095] During media transmission between the LA and the RA 698 and 699, media packets originated by the LA are sent to the LME for optional encryption, then to the LMM, which encapsulates the media packets by appending a tag appropriate for the LA-RA connection and forwards the encapsulated packets over the pre-allocated reservation protocol session 690 to the RMDX. The RMDX determines the RA is the intended destination based upon the tag, removes the tag, and forwards the media packet to the RA via the RMD. Media packets originated by the RA are sent to the RME which performs encryption then to the RMM which encapsulates the media packets by appending a tag appropriate for the LA-RA connection and forwards the encapsulated packets over the pre-allocated reservation protocol session 690 to the LMDX. The LMDX determines the LA is the intended destination based upon the tag, removes the tag, and forwards the media packet to the LA via the local media decryptor (LMD).

An Exemplary H.323 VoIP Implementation

[00096] H.323 is basically an umbrella that covers several existing protocols, including but not limited to H.225.0, and H.245. The later two protocols are used to establish call connection, and capability information between two endpoints. Once this information is exchanged, the endpoints may use RTP and RTCP to exchange voice, voice-band data, and multi-media information between them.

[00097] H.323 suggests that RTP/RTCP should be established between two endpoints (caller/receiver) for each call. Consequently, in order to provide QoS for each call using a protocol like RSVP would mean that every endpoint pair (caller/receiver) for every H.323 call would need to establish RSVP between one another. This would create a huge amount of overhead on the endpoint and adversely affect network resources as RSVP “soft states” must be maintained for the life of the call. This quickly becomes a tremendous scalability issue, since as number of simultaneous calls increase, so do the RSVP “soft state” maintenance messages between endpoints, and every router involved in the transmitting RTP/RTCP data stream.

[00098] Embodiments of the media aggregation manager 300 described herein seek to provide a clean, and scalable solution for this problem, while providing the same QoS as if two individual endpoints had used a reservation protocol session, such as RSVP, between them. Briefly, in the context of a H.323 VoIP embodiment, the H.323 endpoints (callers/receivers) need not have knowledge of how to establish and maintain RSVP sessions. Instead, the media aggregation managers may establish one or more RSVP “pipes” between them that can accommodate several (expected) voice calls. These RSVP pipes are created as the media aggregation managers are started and the RSVP pipes are dynamically maintained between them. This immediately reduces the amount of RSVP state processing in the network. The RSVP pipes between media aggregation managers may be created based upon an educated estimate of the number of calls that are expected between user communities being managed by these media aggregation managers. Since RSVP by nature is established between a specific IP address/port pair and since the pipes are pre-created between media

aggregation managers, all voice traffic (e.g., RTP/RTCP) originates and terminates between media aggregation managers at the media multiplexor 350 and the media demultiplexor 360, respectively.

[00099] In this manner, according to one embodiment, the “local” media aggregation manager appears to an H.323 voice application caller as its intended receiver. The H.323 endpoints make calls to the local media aggregation managers without realizing the local media aggregation managers are not really the final destination. The local media aggregation manager calls the remote media aggregation manager and passes the RTP/RTCP voice data to it. The remote media aggregation manager receives the voice data and sends it the “real” receiver while hiding all multiplexing details from both the caller and the receiver. However, as the voice data is actually exchanged between media aggregation managers over the network it gets RSVP treatment, reserved bandwidth, and QoS. Advantageously, this solution serves as a surrogate to route calls over the pre-created RSVP pipes eliminating QoS processing by endpoints, without any deviations from each involved standard protocol.

[000100] Referring now to **Figure 7**, a flow diagram illustrating exemplary Registration, Admission, Status (RAS) signaling processing will now be described. At decision block 710, the appropriate processing path is determined based upon the triggering event. If the event is a request for a terminal’s signaling address then processing proceeds to decision block 720. If the event represents a signaling address response, then control flow branches to processing block 750. However, if the event is a new call request, then processing continues with decision block 760.

[000101] At decision block 720, in response to a request for a terminal signaling address, a determination is made whether or not the terminal is locally serviced. If it is determined that the terminal is not serviced by the media aggregation manager 300, then processing continues with processing block 730; otherwise processing proceeds to processing block 740.

[000102] At processing block 730, the media aggregation manager 300 requests the call signaling address from an appropriate remote media aggregation manager. For example, the

local media aggregation manager may transmit a multicast message or a directed broadcast to locate the appropriate remote media aggregation manager that services the desired terminal.

[000103] At processing block 740, the media aggregation manager 300 returns its own signaling address rather than the signaling address of the locally serviced terminal. In this manner, subsequent call signaling and control signaling is routed through the local media aggregation manager rather than letting the locally service terminal handle this signaling directly.

[000104] At processing block 750, in response to a signaling address response, the media aggregation manager 300, as above, returns its signaling address in place of the signaling address of the locally serviced terminal to abstract call and control signaling from the locally serviced terminal.

[000105] At decision block 760, in response to a new call request on the RAS channel of the media aggregation manager 300, a determination is made whether there is capacity for the new call. For example, the local resource manager verifies whether the reservation protocol session over which the new call will be multiplexed can accommodate the additional bandwidth requirements of the new call. At any rate, if the local resource manager determines that the reservation protocol session has adequate resources for the new call, then processing continues to processing block 770. Otherwise, control flows to processing block 780.

[000106] At processing block 770, the media aggregation manager 300 returns an indication that the new call can be accepted. At processing block 780, the media aggregation manager 300 returns direction to reject the new call.

[000107] Advantageously, since the terminals/phones register with the media aggregation manager 300, additional authentication processing can be performed in addition to optional encryption, thereby also serving as a checkpoint for only accepting packets from those entities/end points/phones that have previously registered.

[000108] **Figure 8** is a flow diagram illustrating call signaling processing according to one embodiment of the present invention. At decision block 810, the appropriate processing

path is determined based upon the event that has triggered the call signaling processing tread. If the event is a local call connect request, the processing proceeds to processing block 820. If the event represents a remote call connect request, then control flow branches to processing block 830. If the event is a local alerting/call or proceeding/connect message, then processing continues with processing block 840. However, if the event is a remote alerting/call or proceeding/connect message, the processing proceeds with processing block 850.

[000109] At processing block 820, in response to a local call connect request, the media aggregation manager 300 accepts the call from the local terminal and calls the remote media aggregation manager that services the destination terminal. In this manner, the local media aggregation manager poses as the intended receiver to its local terminals that are callers.

[000110] At processing block 830, in response to a remote call connect request, the media aggregation manager 300 accepts the call from the remote media aggregation manager and calls the intended recipient, e.g., one of the terminals serviced by the local media aggregation manager. In this manner, the local media aggregation manager poses a caller to its local terminals that are receivers.

[000111] At processing block 840, in response to a local alerting/call or proceeding/connect message, the local media aggregation manager relays the message to the appropriate remote media aggregation manager(s).

[000112] At processing block 850, in response to a remote alerting/call or proceeding/connect message, the local media aggregation manager relays the message to the appropriate local terminal(s). After processing block 850, call signaling is complete and control protocol signaling (e.g., H.245) can begin.

[000113] **Figure 9** is a flow diagram illustrating control signaling processing according to one embodiment of the present invention. At decision block 910, the appropriate processing path is determined based upon the event that has triggered the control signaling processing tread. If the event is receipt of a master/slave and capability exchange from a local application/endpoint, the processing proceeds to processing block 920. If the event represents receipt of a master/slave and capability exchange from a remote media aggregation manager,

then control flow branches to processing block 930. If the event is receipt of logical channel information from a local application/endpoint, then processing continues with processing block 940. However, if the event is reception of logical channel information from a remote media aggregation manager, the processing proceeds with processing block 950.

[000114] At processing block 920, the master/slave and capability exchange is transmitted to the remote media aggregation manager.

[000115] At processing block 930, the master/slave and capability exchange is transmitted to the local application/endpoint. At processing block 940, the logical channel information from the local application/endpoint is stored in anticipation of making a connection with the media and/or control channels of the local application/endpoint.

[000116] At processing block 950, the LMAM forwards its own logical channel information to the RMAM. Additionally, the network address of the LA is sent to the RMAM.

[000117] At processing block 960, the network address of the RA is stored in a lookup table for address translation and the logical channel information of the LMAM is forwarded to the LA.

[000118] **Figure 10** is a flow diagram illustrating media/control transmission processing according to one embodiment of the present invention. At processing block 1010, the local media multiplexor reports the resources being consumed by the local application/endpoint to the local resource manager.

[000119] At processing block 1020, the media aggregation manager 300 connects to the media and/or control channels of the local application/endpoint.

[000120] At processing block 1030, media and control packets generated by the local application/endpoint are received by the local media encryptor (LME). Depending upon the application session with which the media packets are associated the appropriate form of encryption is applied to the media packets at processing block 1035.

[000121] According to this example, at processing block 1040, after optional encryption is performed, the media multiplexor 350 marks the outbound packets with appropriate

address information (referred to as a “tag”) for demultiplexing at the remote media aggregation manager. The tag is typically appended to transport protocol packets, such as TCP or RTP packets, to allow the media multiplexor 350 to direct packets to the appropriate remote application/endpoint. According to one embodiment, the tag includes address information, such as the destination network address associated with the remote application/endpoint. The destination network address may be determined with reference to a lookup table that allows translation between the source network address associated with the local application/endpoint and the destination network address associated with the remote application/endpoint. Alternatively, a lookup table may be maintained on the media demultiplexor 360 and the tag would include the source network address associated with the local application/endpoint. Then, the source network address would be used by the remote media demultiplexor to determine how to route the inbound packet to the appropriate remote application/endpoint.

[000122] When different channels or ports are used for transport and control protocols (such as RTP and RTCP), then the tag may also include additional protocol dependent control information to allow multiplexing of data and control packets onto the reservation protocol session. Therefore, at optional processing block 1050, each outbound packet may additionally be marked as control or data to allow the remote media aggregation manager to determine the appropriate channel/port of the remote application/endpoint on which to forward the packet.

[000123] Finally, at processing block 1060, the marked packet is transmitted to the appropriate remote media aggregation manager(s).

[000124] **Figure 11** is a flow diagram illustrating media/control reception processing according to one embodiment of the present invention. At processing block 1110, a packet is received from a remote media aggregation manager. The demultiplexing information (e.g., the tag) added by the remote media multiplexor is stripped from the packet and examined at processing block 1120. Optionally, at processing block 1130, if control and data packets are

being multiplexed onto the reservation protocol session, a determination is made whether the packet is a media packet or a control packet based upon the tag. Encrypted media packets are decrypted at processing block 1135 using the appropriate form of decryption for the associated application session. At processing block 1140, the appropriate local application(s)/endpoint(s) to which the packet is destined is/are determined. As described above, the media multiplexor 350 may perform address translation from a source network address to a destination network address. In this case, the appropriate local application(s)/endpoint(s) that are to receive the packet is/are determined by examining the address portion of the tag. Alternatively, if the media multiplexor 350 leaves the source network address in the address portion of the tag, then the appropriate local application(s)/endpoint(s) is/are determined by first translating the address portion using a local lookup table, for example.

[000125] In any event, finally, at processing block 1150, the packet is transmitted to those of the local application(s)/endpoint(s) identified in processing block 1140. If, according to the particular transport and/or control protocols employed, the application(s)/endpoint(s) receive media packets and control packets on different channels/ports, then the packet is forwarded onto the appropriate channel/port of the local application(s)/endpoints(s) based on the packet classification performed at processing block 1130.

[000126] **Figure 12** conceptually illustrates application session establishment in an H.323 environment according to one embodiment of the present invention. In general, the media aggregation managers may abstract the true application session endpoints from each other and serve as proxies for their respective local applications/endpoints. As explained above, the media aggregation managers accomplish this by intercepting signaling messages originating from their respective local applications/endpoints and modifying the signaling messages to make themselves appear as the actual callers/recipients.

[000127] In this illustration, for simplicity, it is assumed that a single local application/endpoint (LA) is establishing an application session with a single remote application/endpoint (RA) over a pre-allocated reservation protocol session 1290 between a

local media aggregation manager (LMAM) geographically proximate to the LA and a remote media aggregation manager (RMAM) geographically proximate to the RA.

[000128] According to this example, application session establishment involves RAS signaling 1210 and 1230, H.225 signaling 1240, and H.245 signaling 1250. RAS signaling 1210 begins with a request for the RA signaling address 1211 by the LA to the LMAM. The LMAM transmits the request 1211 via the reservation protocol session 1290 to the RMAM. In response to the request 1211, the RMAM decides it wants to route H.225/H.245 signaling through it instead of letting the RA do it directly. Therefore, the RMAM replies with a packet 1212 containing RMAM's signaling address. Similarly, the LMAM decides it wants to route H.225/H.245 signaling through it instead of letting the LA do it directly. Therefore, the LMAM substitutes its signaling address for that of the RMAM and forwards packet 1213 to the LA.

[000129] RAS signaling continues with the RA asking the RMAM (on its RAS channel) if it is okay to accept a new call by sending the RMAM a new call request 1231. The RMAM authorizes the new call by responding with a packet 1231 giving the RA permission to accept the new call.

[000130] H.225 signaling comprises the RA sending H.225 alerting/call proceeding/connect messages 1241 to the RMAM. The RMAM sends the same to the LMAM; and the LMAM sends the same to the LA. At this point, the LA determines that H.225 call signaling is complete and starts H.245 signaling.

[000131] H.245 signaling begins with the LA sending master/slave and capability exchange messages 1251 to the LMAM, which are relayed to the RMAM and from the RMAM to the RA. Then, the RA sends master/slave and capability exchange messages 1252 to the RMAM. The RMAM transmits these messages to the LMAM; and the LMAM forwards them to the LA.

[000132] Subsequently, the LA initiates an exchange of logical channel information by sending logical channel information packets 1253 to the LMAM. The logical channel information identifies the network address (e.g., IP address) and port numbers where

RTP/RTCP connections will be accepted. The LMAM stores the LA's logical channel information and passes its own connection information 1254 to the RMAM. Additionally, the LMAM provides the network address of the LA to the RMAM for later use in address translation lookups. As mentioned above, the network address of the LA may be used by the RMM or the RMDX depending upon where the address translation lookup is performed. The RMAM remembers the information provided by the LMAM and generates its own RTP/RTCP information 1255 and passes it to the RA.

[000133] After receiving logical channel information thought to be associated with the LA, the RA sends its logical channel information 1256 to the RMAM (thinking it is being directed to the LA). The RMAM stores the RA's logical channel information and passes its own connection information 1257 to the LMAM. Additionally, the RMAM provides the network address of the RA to the LMAM. The LMAM remembers the logical channel information provided by the RMAM and generates its own RTP/RTCP information 1258 and passes it to the LA.

[000134] The LA sends an ACK message 1259 to the LMAM to acknowledge receipt of what it thinks to be the RA's logical channel information. The acknowledgement is relayed to the RA by the LMAM and the RMAM. The RA also sends an ACK message 1260 to the RMAM to acknowledge receipt of what it thinks to be the LA's logical channel information. The acknowledgement is related to the LA by the RMAM and the LMAM. Finally, the LMAM and the RMAM each use the logical channel information intercepted from the LA and the RA, respectively, to connect to the media and/or control channels of the LA and RA.

Exemplary Bridging Between Potentially Heterogeneous IP Telephony Environments

[000135] **Figure 13** conceptually illustrates H.323 signaling and media flow according to an embodiment in which call management is performed external to the media aggregation managers. Various embodiments described earlier were discussed as if call management signaling was performed by the media aggregation manager 300. According to the example illustrated by Figure 13, call management software running on or interacting with one or more

IP PBXs performs call management signaling and the media aggregation manager 300 performs signaling gateway and media gateway functionality and acts as an H.323 gateway for the IP PBX(s).

[000136] As mentioned earlier, vendors of hardware and software for IP telephony products typically employ proprietary signaling protocols (e.g., 1310 and 1320) thereby requiring customers to implement homogeneous VoIP environments. According to embodiments of the present invention, however, customers are provided with the flexibility to implement heterogeneous VoIP environments because the media aggregation manager 300 translates among various proprietary vendor signaling protocols thereby allowing IP PBX systems and IP phone sets from different vendors to interoperate.

[000137] In this signaling and media flow example, external H.323 signaling is performed by the IP PBX call management agent (e.g., CM 130 and CM 150) on behalf of the phone endpoints. At 1330, the call management agents 130 and 150 register with the main media aggregation manager gatekeeper. At 1335, IP phone 111 initiates a call to IP phone 121 by sending a setup message to the CM 130. In response, the CM 130 sends an admission request (ARQ) message to the main media aggregation manager which confirms the admission (ACF), 1340 . The main media aggregation manager knows that media aggregation manager 115 is the local media aggregation manager for IP phone 111 and returns the address of media aggregation manager 115 in the ACF response message to the CM 130. The CM 130 now knows to send future signaling messages to media aggregation manager 115.

[000138] At 1345, the CM 130 sends the setup message to the media aggregation manager 115. In response, at 1350, the media aggregation manager 115 issues a location request message (LRQ) to resolve the media aggregation manager for the destination IP phone 121. Media aggregation manager 125 is the local owner for IP phone 121 and at 1355 it sends the location confirm (LCF) to media aggregation manager 115.

[000139] At 1360, media aggregation manager 115 sends the setup message to media aggregation manager 125 and media aggregation manager 125 forwards the setup message to

CM 150 (Note: a media aggregation manager knows its associated call management agent via configuration).

[000140] At 1365, CM 150 notifies IP phone 121 of the call and sends the call proceeding message to media aggregation manager 125. CM 150 sends an admission request message (ARQ) for IP phone 121 to the main media aggregation manager, 1370, and the main media aggregation manager confirms the admission (ACF), at 1375.

[000141] At 1380, the alerting/connect messaging proceeds as usual and at 1390 H.245 open logical channel messaging takes place to open the media and signaling channels between IP phone 111 and IP phone 121. In the open logical channel acknowledgement messages, the media aggregation managers replace the IP phone's RTP/RTCP IP port pair with its own local addresses which force media to be sent to the media aggregation managers and not directly to the destination phone endpoint. Thereafter, the media aggregation managers 115 and 125 ensure media packets associated with the media/control stream 171 are forwarded over the predefined IPVC QoS pipe 170.

Exemplary Encapsulated Packet Formats

[000142] **Figure 14A** illustrates the encapsulated ("MUX") packet format 1400 according to one embodiment of the present invention in which address replacement is performed by the LMAM. The payload of the encapsulated packet 1400 includes a destination network address field 1410, a variable length transport or control protocol packet portion 1415, and a packet type indication 1420. The destination network address 1410 is typically the IP address of the true recipient (e.g., the application/endpoint to which the packet is destined). In environments where multiplexing of control and data is employed, the variable length portion 1415 may include either a transport protocol packet (e.g., a RTP packet) or a control protocol packet (e.g., a RTCP packet) as indicated by the packet type indication 1420. In alternative embodiments, where multiplexing of control and data is not employed, then the variable length portion 1415 would still include either control or data, but the packet type indication 1420 would no longer be necessary.

[000143] **Figure 14B** illustrates media transmission in both directions according to the encapsulated packet format of **Figure 14A**. When the LA originates a media packet, it generates a packet 1440 including media 1442. The LMAM optionally encrypts the media 1442 and encapsulates the media 1442 in the encapsulated packet format 1400 by generating an encapsulated packet 1450 that includes the RA's network address 1451, the media 1442, and a packet type indicator 1453. For example, upon receipt of packet 1440, the LMAM may append the network address of the RA and a packet type indicator 1453 based upon the channel/port upon which the packet 1440 was received. When the encapsulated packet 1450 is received by the RMAM, it strips the information added by the LMAM, decrypts the media 1442, if necessary, and forwards a packet 1460 comprising the media 1442 to the RA.

[000144] When the RA originates a media packet, it generates a packet 1490 including media 1492. The RMAM optionally encrypts the media 1492 and encapsulates the media 1492 in the encapsulated packet format 1400 by generating an encapsulated packet 1480 that includes the LA's network address 1441, the media 1492, and a packet type indicator 1483. For example, upon receipt of packet 1490, the RMAM may append the network address of the LA and a packet type indicator 1483 based upon the channel/port upon which the packet 1490 was received. When the encapsulated packet 1480 is received by the LMAM, it strips the information added by the RMAM, decrypts the media 1492, if necessary, and forwards a packet 1470 comprising the media 1492 to the LA.

[000145] **Figure 15A** illustrates the encapsulated ("MUX") packet format according to another embodiment of the present invention in which address replacement is performed by the RMAM. The payload of the encapsulated packet 1500 includes a source network address field 1510, a variable length transport or control protocol packet portion 1515, and a packet type indication 1520. The source network address 1510 is typically the IP address of the true caller (e.g., the application/endpoint from which the packet is originated). In environments where multiplexing of control and data is employed, the variable length portion 1515 may include either a transport protocol packet (e.g., a RTP packet) or a control protocol packet (e.g., a RTCP packet) as indicated by the packet type indication 1520. In alternative

embodiments, where multiplexing of control and data is not employed, then the variable length portion 1515 would still include either control or data, but the packet type indication 1520 would no longer be necessary.

[000146] **Figure 15B** illustrates media transmission in both directions according to the encapsulated packet format of **Figure 15A**. When the LA originates a media packet, it generates a packet 1540 including media 1542. The LMAM optionally encrypts the media 1542 and encapsulates the media 1542 in the encapsulated packet format 1500 by generating an encapsulated packet 1550 that includes the LA's network address 1541, the media 1542, and a packet type indicator 1553. For example, upon receipt of packet 1540, the LMAM may append the network address of the LA and a packet type indicator 1553 based upon the channel/port upon which the packet 1540 was received. When the encapsulated packet 1550 is received by the RMAM, it strips the information added by the LMAM, decrypts the media 1542, if necessary, and forwards a packet 1560 comprising the media 1542 to the RA by looking up the network address of the RA based upon the LA's network address 1541.

[000147] When the RA originates a media packet, it generates a packet 1590 including media 1592. The RMAM optionally encrypts the media 1592 and encapsulates the media 1592 in the encapsulated packet format 1500 by generating an encapsulated packet 1580 that includes the RA's network address 1551, the media 1592, and a packet type indicator 1583. For example, upon receipt of packet 1580, the RMAM may append the network address of the RA and a packet type indicator 1583 based upon the channel/port upon which the packet 1580 was received. When the encapsulated packet 1580 is received by the LMAM, it strips the information added by the RMAM, decrypts the media 1592, if necessary, and forwards a packet 1570 comprising the media 1592 to the RA by looking up the network address of the LA based upon the RA's network address 1551.

[000148] In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the

invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.
